RIP-2 (1993)

- Extensions to original RIP attempt to address RIP's problems
- Authentication with a cleartext password
- Passes subnet mask
- Next-hop IP address allows interaction with OSPF and EGPs.

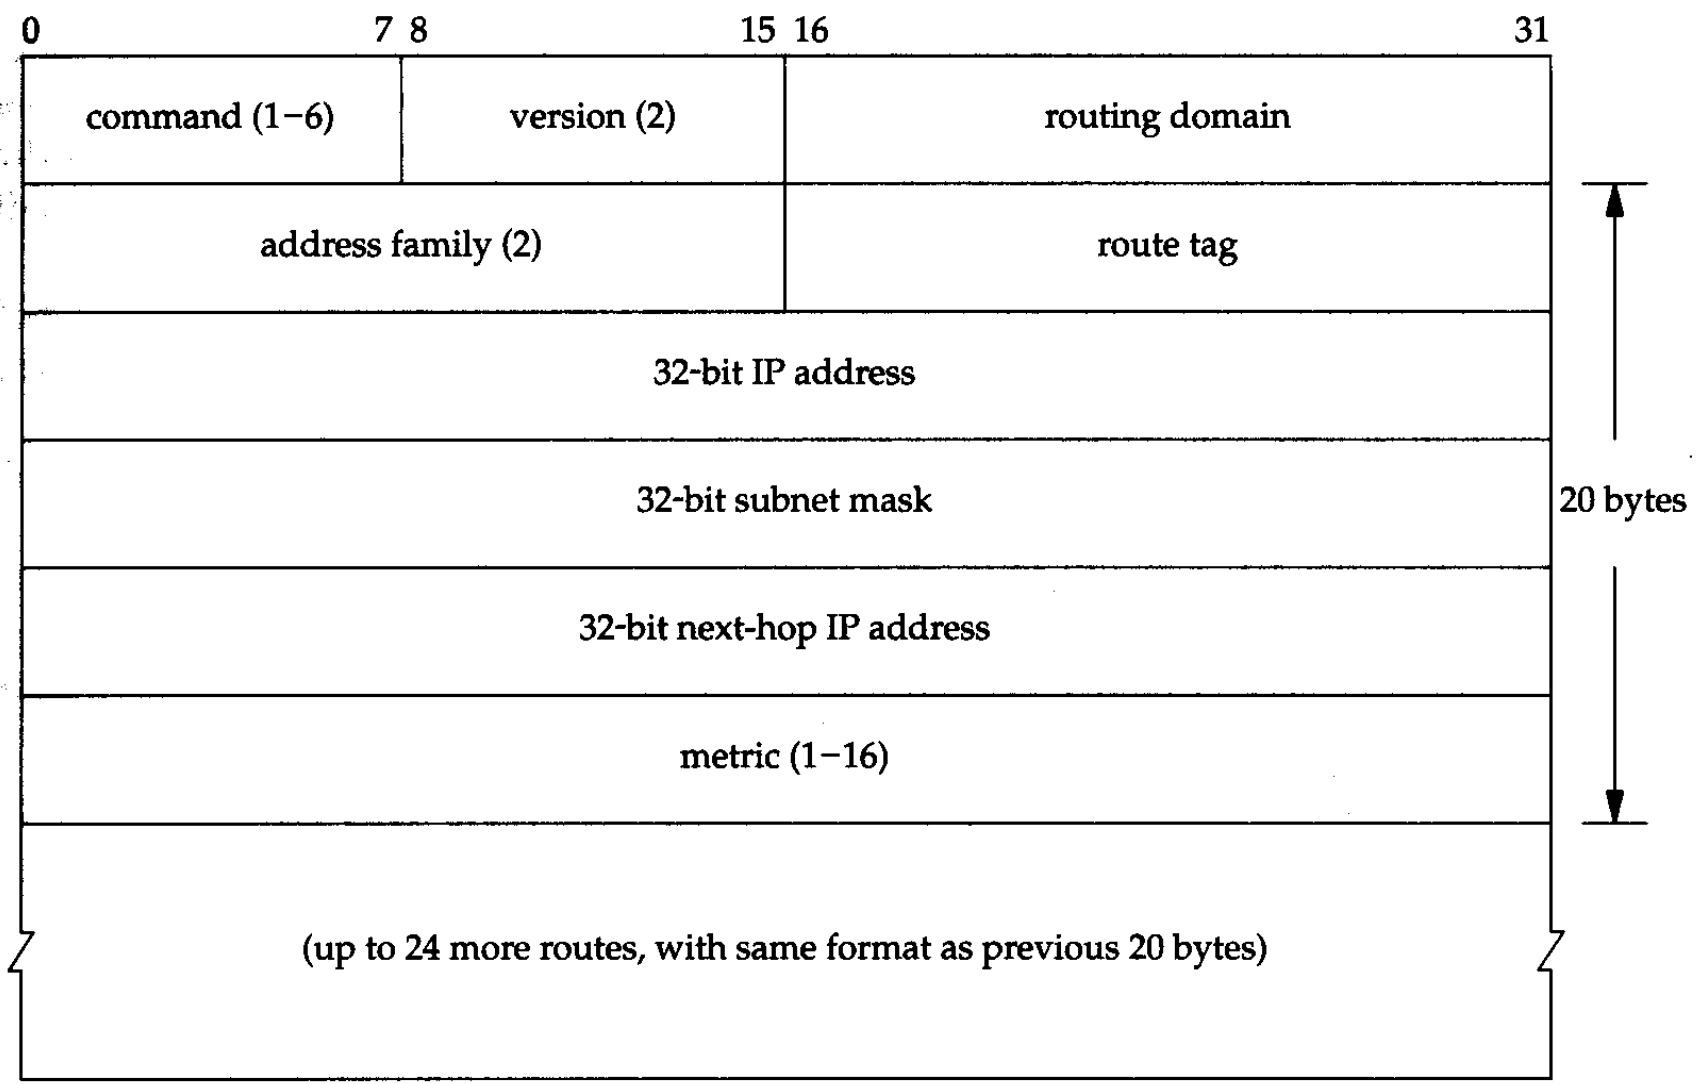
0
7 8
15 16
31
command (1–6)
version (2)
routing domain
address family (2)
route tag
32-bit IP address
32-bit subnet mask
32-bit next-hop IP address
metric (1–16)
20 bytes
(up to 24 more routes, with same format as previous 20 bytes)

Open Shortest Path First (OSPF)

- Link-state protocol - Djikstra's algorithm
- A single "link-state" table is maintained by all routers in an AS
- Each router is responsible for maintaining and disseminating information about it's interfaces in the "link-state" table

OSPF (cont)

- A router passes “link-states” to all other routers using flooding
- Remaining routers enter this information into their copy of the “link-state” table
- Each router then computes its own routing table from the values in the link state table

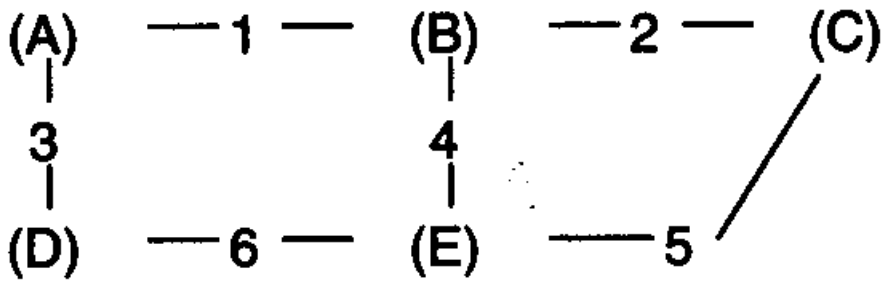

From	To	Link	Distance
A	B	1	1
A	D	3	1
B	A	1	1
B	C	2	1
B	E	4	1
C	B	2	1
C	E	5	1
D	A	3	1
D	E	6	1
E	B	4	1
E	C	5	1
E	D	6	1

Typical Flooding Message from A

From A, to B, link 1
distance = 1
From A, to D, link 3
distance = 1

Why is OSPF Better?

- Fast, loopless convergence
  - Given N nodes and M links, Bellman-Ford converges in O(NM), Djikstra converges in O(MlogM)
  - RIP uses distributed computation - number of steps required is proportional to the number of nodes
  - OSPF floods then does a local computation

Why is OSPF Better? (cont)

- Support multiple metrics simultaneously (corresponding to IP TOS)
  - Maximum Throughput
  - Minimize Delay
  - Minimize Monetary Cost
  - Maximize Reliability
- Can calculate a separate set of routes for each

Why is OSPF Better? (cont)

- Supports multiple paths to a destination
  - Availability of “link-state” table allows determination of alternative routes
  - Can distribute traffic between routes to minimize congestion => “load balancing”
- Support representation of external routes

OSPF Design

- Protocol runs directly on top of IP (protocol type 89)
- Composed of three subprotocols: hello, exchange, and flooding
- Additionally, five message types: Router Links, Network Links, Summary Links, Summary to a border, and External Links

Exterior Routing Protocols

- Exterior Gateway Protocol (EGP)
  - First interdomain routing protocol
  - Viewed the Internet as a tree topology with a single backbone (NSFNet)
  - Autonomous Systems (AS) connected only as parents and children and not as peers

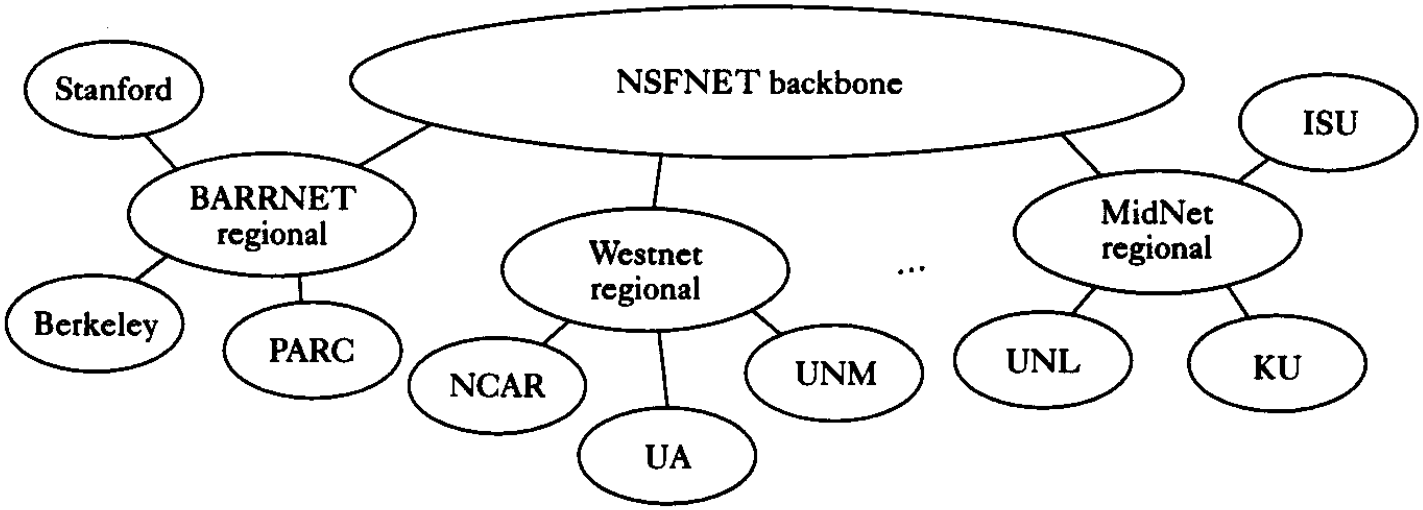

Figure 5.12 The tree structure of the Internet in 1990.

Border Gateway Protocol (BGP)

- Assumes the Internet is an “arbitrarily connected” interconnected set of ASs
- Currently in its fourth version (BGP-4)
- Implemented within “gated”
- Facilitates multiple-backbone “service provider networks” owned by private companies

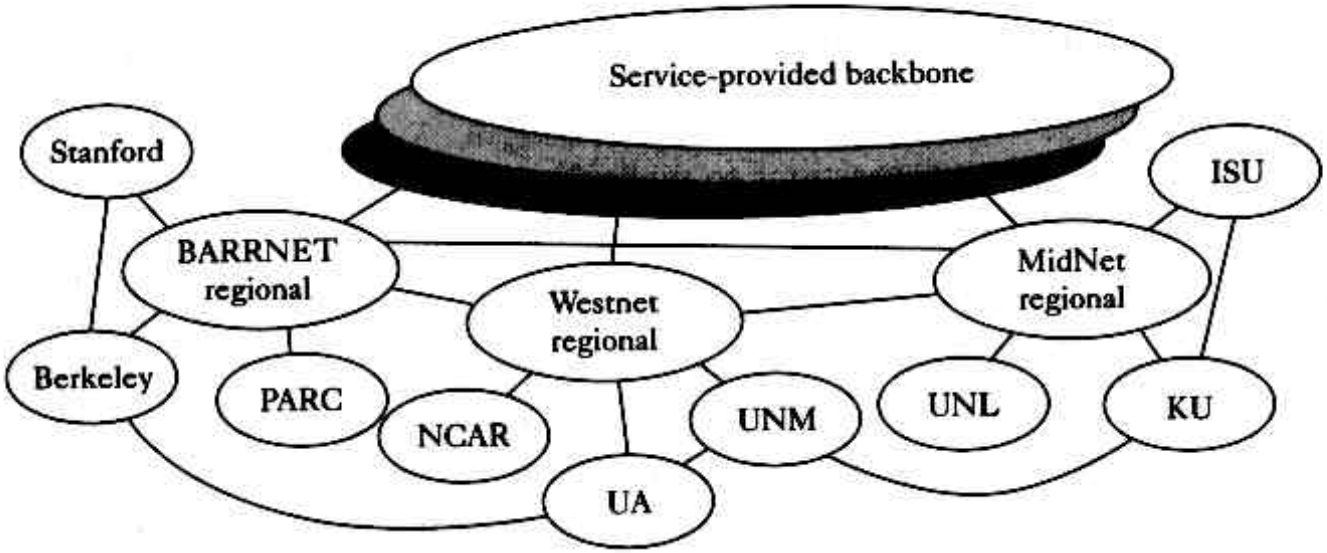

Figure 5.18 Today's multibackbone Internet.

More about ASs

- “Local traffic” is defined as traffic that originates at or terminates on nodes within an AS
- “Transit traffic” is traffic that passes through an AS

AS Types

- “stub AS” - has a single connection to one other AS; only carries local traffic
- “multihomed AS” - has connections to more than one other AS but refuses to carry transit traffic
- “transit AS” - has connections to more than one other AS and is designed to carry both transit and local traffic

BGP Operations

- Each AS has at least one “BGP Speaker”
- In addition, the AS has one or more border gateways, which need not be the same as speakers
- Border gateways act as interfaces to other ASs

BGP Operations (cont)

- Primary goal: find any path to the intended destination that is loop free
- BGP is more concerned about reachability than optimality
- BGP speaker advertises reachability information for all networks within its AS
- In the case of transit ASs, speaker also advertises networks that can be reached through the AS

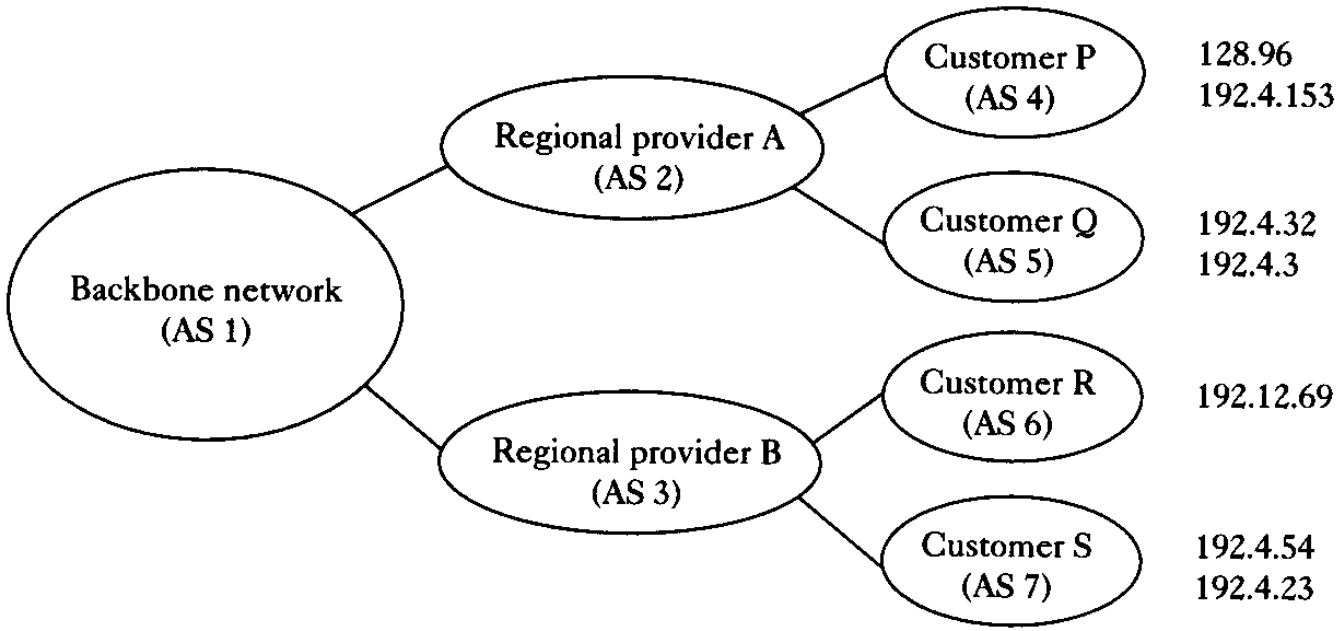

Figure 5.19 Example of a network running BGP.

BGP Operations (cont)

- Establishes TCP connection (Port 179) for information exchange
- Initially, two systems exchange entire BGP routing tables
- Incremental updates are sent as routing table changes (or unless connection is broken)

BGP Operations (cont)

- Detects failure of a link by sending a “keepalive” message (different than TCP “keepalive”)
- Messages sent every 30 seconds.
- If link goes down, a negative advertisement known as a “withdrawn route” is sent

BGP Operations (cont)

- BGP does not belong to either of the two main classes of routing protocols
- BGP advertises complete paths as an enumerated list of ASs to reach a particular network
- Necessary to enable policy decisions of certain ASs
- Also allows detection of routing loops

Acceptable Use Policy (AUP)

- Originally dictated by NSF - only ASs devoted to science could transit
- Led to the development of commercial providers allowing transit
- In turn, led to the idea of “peering” between service providers

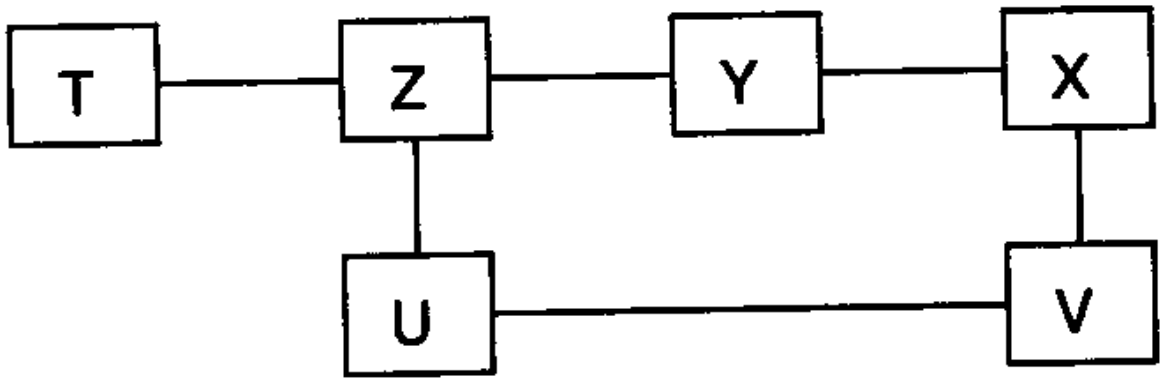

The connectivity for AS Z

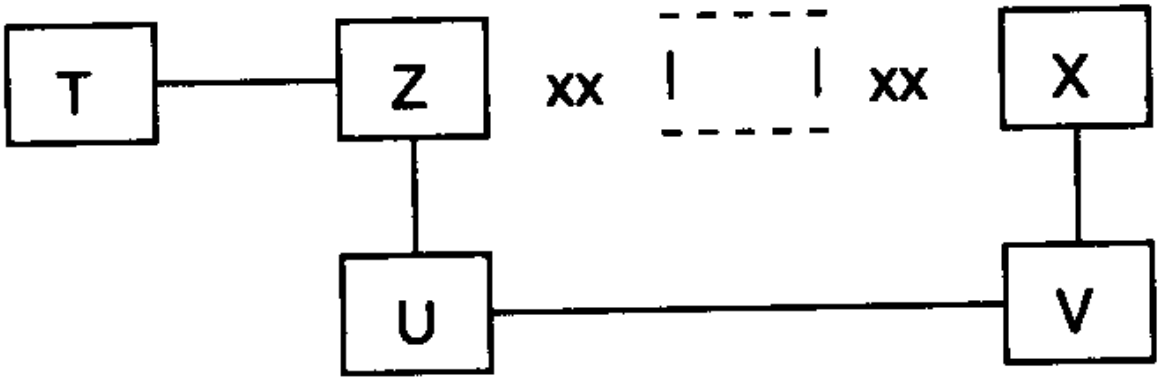

The connectivity for AS T:
transit through Y is banned

BGP Resources

- AS numbers are assigned by either the American Registry for Internet Numbers (ARIN, www.arin.net, Americas, Carib, Africa), RIPE-NCC (Europe), or AP-NIC (Asia).
- http://www.iana.org/numbers.html
- ftp://rs.arin.net/netinfo/asn.txt (Current ASs)